

Decomposition of Multi-Agent Planning under Distributed Motion and Task LTL Specifications

Jana Tumova and Dimos V. Dimarogonas

Abstract—The aim of this work is to introduce an efficient procedure for discrete multi-agent planning under local complex temporal logic behavior specifications. While the first part of an agent’s behavior specification constraints the agent’s trace and is independent, the second part of the specification expresses the agent’s tasks in terms of the services to be provided along the trace and may impose requests for the other agents’ collaborations. To fight the extreme computational complexity of centralized multi-agent planning, we propose a two-phase automata-based solution, where we systematically decouple the planning procedure for the two types of specifications. At first, we only consider the former specifications in a fully decentralized way and we compactly represent each agents’ admissible traces by abstracting away the states that are insignificant for the satisfaction of their latter specifications. Second, the synchronized planning procedure uses only the compact representations. The satisfaction of the overall specification is guaranteed by construction for each agent. An illustrative example demonstrating the practical benefits of the solution is included.

I. INTRODUCTION

Automated synthesis of correct-by-design controllers complying with complex behavior specifications has recently found a wide support in formal verification methods and temporal logics. A number of the suggested solutions build on a hierarchical approach [2], [9], [11], [17], where the dynamics of the system is abstracted into a finite, discrete transition system, a discrete plan that meets the specification is synthesized and next translated into a controller for the original system. This work focuses on a multi-agent version of the above problem, and particularly on the second step of the hierarchical approach. We consider a heterogeneous team of agents (e.g., robots), that are assigned a behavior specification each, comprising of the requirements on the sequence of states they go through (e.g., the robots’ trajectories) and the services they provide (e.g., loading an object). The agents’ tasks are though not completely independent; in order to achieve their task a collaboration of the other agents might be required (e.g., helping to load a heavy object). Our aim in this paper is to efficiently synthesize a strategy for each agent, such that all requirements are met. As a specification language, we use Linear Temporal Logic (LTL), for its resemblance to natural language [7], and expressive power.

This work was supported by EU STREP RECONFIG, European Unions Horizon 2020 Research and Innovation Programme under the Grant Agreement No.644128 (AEROWORKS), Swedish Research Council (VR), and H2020 European Research Council (ERC) Starting Grant BUCOPHSYS. The authors are with the ACCESS Linnaeus Center, School of Electrical Engineering, KTH Royal Institute of Technology, Stockholm, Sweden and with the KTH Centre for Autonomous Systems. Email: tumova@kth.se, dimos@kth.se.

The goal is viewed from the bottom-up perspective, where we introduce the notion of local specification satisfaction based on the individual agents’ own viewpoints.

Multi-agent planning from temporal logic specification has been explored in several recent works. Planning for a team of robots from a global LTL specification was addressed, e.g., in [10], [12]. The authors in [8] considered LTL for vehicle routing and in [16] for search and rescue missions. A decentralized control of a robotic team from local LTL specification with communication constraints is proposed in [4], [6]. In contrast to our work, the agents therein do not impose any constraints on the other agents’ behavior. In [3], [15], a top-down approach to LTL planning is investigated; the team is given a global specification and an effort is made to decompose the formula into independent local specifications that can be treated separately. In [5], bottom-up planning from LTL specifications is considered, and a partially decentralized solution is proposed that decomposes the group into clusters of dependent agents. To cope with the state space explosion, in [13], [14] we have proposed a receding horizon approach to multi-agent planning.

In this paper, we take into account an extension of the previous problem, where besides the high-level *tasks* in terms of services to be provided, the agents’ behaviors are limited by additional temporal logic constraints, allowing to express safety, surveillance, sequencing, or reachability properties of their traces, which we call *motion* specifications for simplicity. We propose to decouple the planning for the former and the latter parts of the specification. The main idea is to build local synchronized product automata for each of the agents’ motion separately and reduce their size before creating the global product. The contribution of this paper can thus be summarized as the introduction of an efficient, bottom-up control strategy synthesis for multi-agent systems from local LTL motion and task specifications. The designed strategy allows the agents to execute their finite plans to a large extent independently, in an asynchronous manner.

The rest of the paper is structured as follows. In Sec. II, we fix necessary notation. Sec. III formalizes the problem statement. In Sec. IV, the details of the solutions are provided. An illustrative example is given in Sec. V.

II. NOTATION AND PRELIMINARIES

Given a set S , let 2^S , and S^ω denote the set of all subsets of S , and the set of all infinite sequences of elements of S , respectively. The i -th element of the sequence w is $w(i)$, and a *subsequence* of an infinite sequence $w = w(1)w(2)\dots$ is a finite or infinite sequence $w(\iota_1)w(\iota_2)\dots$, such that $1 \leq$

$\iota_j \leq \iota_{j+1}$, for all $1 \leq j$. Boolean operator XOR (exclusive or) is denoted by \oplus , positive integers and non-negative real numbers are denoted by \mathbb{N} and \mathbb{R}_0^+ , respectively.

A *labeled transition system (TS)* is a tuple $\mathcal{T} = (S, s_{init}, A, T, \Pi, L)$, where S is a finite set of states; $s_{init} \in S$ is the initial state; A is a finite set of actions; $T \subseteq S \times A \rightarrow S$ is a partial deterministic transition function; Π is a set of atomic propositions; and $L : S \rightarrow 2^\Pi$ is a labeling function. A *trace* of \mathcal{T} is an infinite alternating sequence of states and actions $\tau = s_1 \alpha_1 s_2 \alpha_2 \dots$, such that $s_1 = s_{init}$, and for all $i \geq 1$, $T(s_i, \alpha_i) = s_{i+1}$.

An *LTL formula* φ over the set of atomic propositions Σ is defined inductively: (i) $\pi \in \Sigma$ is a formula, and (ii) if φ_1 and φ_2 are formulas, then $\varphi_1 \vee \varphi_2$, $\neg \varphi_1$, $X \varphi_1$, $\varphi_1 U \varphi_2$, $F \varphi_1$, and $G \varphi_1$ are each a formula, where \neg and \vee are standard Boolean connectives, and X (next), U (until), F (eventually), and G (always) are temporal operators. The semantics of LTL are defined over infinite words over 2^Π (see, e.g., [1]). A trace $\tau = s_1 \alpha_1 s_2 \alpha_2 \dots$ of \mathcal{T} satisfies an LTL formula φ over Π ($\tau \models \varphi$) iff $L(s_1)L(s_2)\dots$ satisfies φ ($L(s_1)L(s_2)\dots \models \varphi$).

A *Büchi automaton (BA)* is a tuple $\mathcal{A} = (Q, q_{init}, \Sigma, \delta, F)$, where Q is a finite set of states; $q_{init} \in Q$ is the initial state; Σ is an input alphabet; $\delta \subseteq Q \times \Sigma \times Q$ is a non-deterministic transition relation; $F \subseteq Q$ is an accepting condition. A run of the BA \mathcal{B} from a state $q_1 \in Q$ over a word $w = \sigma_1 \sigma_2 \dots \in \Sigma^\omega$ is a sequence $\rho = q_1 q_2 \dots$, such that $(q_i, \sigma_i, q_{i+1}) \in \delta$, for all $i \geq 1$ and it is *accepting* if it intersects F infinitely many times. A word w is *accepted* by \mathcal{B} if there exists an accepting run over w from the state q_{init} . Any LTL formula φ over Π can be algorithmically translated into an equivalent BA \mathcal{B} with $\Sigma = 2^\Pi$. Any BA can be also viewed as a graph and the terminology from graph theory applies.

III. PROBLEM FORMULATION

Two general viewpoints can be taken in multi-agent planning: either the system acts as a team with a common goal, or the agents have their individual, local specifications. Here, we adopt the second viewpoint and we do not look at the global team behavior. In contrast, we propose local specification satisfaction to determine whether an agent's task is fulfilled from its own perspective. Each agent's temporal specification comprises of a *motion* and a *task* formula. Loosely speaking, the motion formula can be viewed as the agent's safety, reachability, surveillance or sequencing, while the task one represents its effectiveness, i.e., the motion formula limits the sequence of states the agent's goes through, while the task formula corresponds to an accomplishment of a high-level task composed of simple services that the agent has the ability to provide. Unlike the motion one, the task formula may involve collaboration requirements.

Our goal is to synthesize a control strategy for each agent that involves (i) the actions to be executed and (ii) the requests for synchronization imposed on other agents. Together, the strategies have to ensure the local satisfaction of each of the motion and task specifications regardless of the time duration of the planned action executions.

A. System Model

Let us consider N agents (e.g., robots in a partitioned environment). Each agent $i \in \mathcal{N}$, where $\mathcal{N} = \{1, \dots, N\}$ has two different types of capabilities: the ability to execute an action primitive, and to synchronize with the others.

1) *Action Execution Capabilities*: The agent i 's action-execution capabilities are modeled as a finite TS $\mathcal{T}_i = (S_i, s_{init,i}, A_i, T_i, \Pi_i, L_i)$. The states of the TS correspond to states of the agents (e.g., the location of the robots in the regions of the environment). The atomic propositions represent inherent properties of the system states (e.g., the robot is in a safe region). The actions represent abstractions of the agent's low-level controllers, and the transitions between the states correspond to the agent's capabilities to execute the actions (e.g., the ability of the robots to move between two regions of the environment). The traces are, roughly speaking, the abstractions of the agents' long-term behaviors (e.g., the robots' trajectories). Each of the agents' action executions takes a certain amount of time. Given a trace $\tau = s_1 \alpha_1 s_2 \alpha_2 \dots$ of \mathcal{T}_i , we denote by $\Delta_{\alpha_j} \in \mathbb{R}_0^+$ the time *duration* of the transition $s_j \xrightarrow{\alpha_j} s_{j+1}$. Note that a transition duration is arbitrary and unknown prior its execution, and that the execution of the same action α may induce different transition durations in its different instances on the trace.

2) *Synchronization Capabilities*: Next to the action-execution capabilities, the agents have the ability to *synchronize*, i.e., to wait for each other and to proceed with the further execution simultaneously. The synchronization is modeled through the *synchronization requests*. While being in a state s , an agent i can send a request $sync_i(I)$ to a subset of agents $\{i\} \subseteq I \subseteq \mathcal{N}$ notifying that it is ready to synchronize. Then, before proceeding with the execution of any action $\alpha \in A_i$, it has to wait till $sync_{i'}(I)$ has been sent by each agent $i' \in I$, i.e., till the moment when each agent $i' \in I$ is ready to synchronize, too. For simplicity, we assume the perfect propagation of the synchronization requests. The synchronization is immediate once each of the agents $i' \in I$ has sent its request $sync_{i'}(I)$ and all agents in I start executing the next action collaboratively, at the same time. The set of all synchronization requests of an agent i is $Sync_i = \{sync_i(I) \mid \{i\} \subseteq I \subseteq \mathcal{N}\}$. Note that the synchronization request $sync_i(\{i\})$ indicates that no synchronization with the others is needed, and the next action of agent i is executed immediately, independently on the other agents. We assume that each agent sends a synchronization request instantly once it completes an action execution and that it starts executing an action instantly once it synchronizes with the other agents. Instead of allowing for idling, we include an existence of a special action $stay_i \in A_i$ and a so-called *self-loop* $s \xrightarrow{stay_i} s$ for all $s \in S_i$, and all \mathcal{T}_i , where $i \in \mathcal{N}$. Given a trace $\tau_i = s_{i,1} \alpha_{i,1} s_{i,2} \alpha_{i,2} \dots$ of \mathcal{T}_i , we denote by $\Delta_{s_{i,j}} \in \mathbb{R}_0^+$ the time duration of the synchronization that has been requested in the state $s_{i,j}$. Note that $\Delta_{s_{i,j}} = 0$ for $sync_i(\{i\})$ in $s_{i,j}$.

3) *Services*: Each of the agents' specification is given via two components: the first one are temporal requirements on

the atomic propositions that need to hold along its trace, and the second one is a task given in terms of events of interest associated with their actions, which we call *services* (e.g., an object pick-up). The set of services that can be *provided* by an agent $i \in \mathcal{N}$ is Σ_i . Services are provided within the agents' transitions; each action $\alpha \in A_i$ is associated, or in other words *labeled*, either with (i) a service set $\sigma \in 2^{\Sigma_i}$ provided by i upon the execution of α , or (ii) a special *silent* service set $\mathcal{E}_i = \{\varepsilon_i\}$, where $\varepsilon_i \notin \Sigma_i$ indicating that the action α does not provide any event of interest. Hence, two additional components of the agent i 's model are the set of all available services Σ_i and the action-labeling function $\mathcal{L}_i : A_i \rightarrow 2^{\Pi_i} \cup 2^{\mathcal{E}_i}$. Note that we specifically distinguish between a silent service set \mathcal{E}_i and an empty service set $\{\}$. The above mentioned self-loops of form $s \xrightarrow{\text{stay}_i} s$ that we use to model the agents' ability of staying in their respective current states are labeled with the silent service set $\mathcal{L}_i(\text{stay}_i) = \mathcal{E}_i$. Without loss of generality, we assume that $\Sigma_i \cap \Sigma_{i'} = \emptyset$, for all $i, i' \in \mathcal{N}$, such that $i \neq i'$, which is necessary for the final step of our solution. Finally, we model an agent $i \in \mathcal{N}$ as the tuple $\mathcal{M}_i = (\mathcal{T}_i = (S_i, s_{\text{init},i}, A_i, T_i, \Pi_i, L_i), \text{Sync}_i, \Sigma_i, \mathcal{L}_i)$.

4) *Behaviors*: Behavior of an agent i is defined by its states, actions, synchronizations with the others, and the time instants of the action executions and the synchronizations.

Definition 1 (Behavior and strategy). A behavior of an agent i is a tuple $\beta = (\tau, \gamma, \mathbb{T})$, where

- $\tau = s_1\alpha_1s_2\alpha_2\dots$ is a trace of \mathcal{T}_i ;
- $\gamma = r_1r_2\dots$ is a synchronization sequence, where $r_j \in \text{Sync}_i$ is the synchronization request sent at s_j ; and
- $\mathbb{T} = t_{s_1}t_{\alpha_1}t_{s_2}t_{\alpha_2}\dots$ is a non-decreasing behavior time sequence, where t_{s_j} is the time instant when the synchronization request r_j was sent, and t_{α_j} is the time instant when the action α_j started being executed. The following properties hold: $t_{s_1} = 0$, and for all $j \geq 1$, $t_{s_{j+1}} - t_{\alpha_j} = \Delta_{\alpha_j}$, and $t_{\alpha_j} - t_{s_j} = \Delta_{s_j}$.

A strategy (τ, γ) for an agent i is a trace τ and a synchronization sequence γ .

The notion of behavior does not reflect the above described inter-agent synchronization rules. To that end, we define behaviors induced by strategies admissible by a team of agents. We denote the behavior of an agent $i \in \mathcal{N}$ by $\beta_i = (\tau_i, \gamma_i, \mathbb{T}_i)$, and we use $\tau_i = s_{i,1}\alpha_{i,1}s_{i,2}\alpha_{i,2}\dots$, $\gamma_i = r_{i,1}r_{i,2}\dots$, and $\mathbb{T}_i = t_{s_{i,1}}t_{\alpha_{i,1}}t_{s_{i,2}}t_{\alpha_{i,2}}\dots$ to denote its trace, synchronization and time sequence.

Definition 2 (Induced behaviors). The set of behaviors induced by a collection of strategies $(\tau_1, \gamma_1), \dots, (\tau_N, \gamma_N)$ of agents in \mathcal{N} are the subset of the collections of their behaviors $\mathbb{B} \subseteq \{\beta_i \in \{\beta_1, \dots, \beta_N \mid \beta_i \text{ is a behavior of agent } i\}\}$ satisfying the following condition for all $i \in \mathcal{N}$, and $j \geq 1$: Suppose that $r_{i,j} = \text{sync}_i(I)$. Then for all $i' \in I$ there exists a matching index $j' \geq 1$, such that $r_{i',j'} = \text{sync}_{i'}(I)$, and $t_{\alpha_{i,j}} = t_{\alpha_{i',j'}}$. Furthermore, there exists at least one $i' \in I$, such that $t_{s_{i',j'}} = t_{\alpha_{i',j'}}$, i.e., such that $\Delta_{s_{i',j'}} = 0$, for the matching index j' .

B. Motion and Task Specifications

The individual agents' tasks are collaborative; on the one hand, they involve requirements on the agents' states, i.e., the atomic propositions associated with them and on the other hand, they concern the respective agent's services as well as the services of the other agents. Formally, each of the agents is given (i) an LTL $_{\setminus X}$ formula ϕ_i over Π_i , which we call a *motion* specification and (ii) a collaborative LTL formula ψ_i over $\Sigma = \bigcup_{i' \in \mathcal{N}} \Sigma_{i'}$, which we call a *task* specification. Consider for a moment a single agent $\mathcal{M}_i = (\mathcal{T}_i, \text{Sync}_i, \Sigma_i, \mathcal{L}_i)$, and its behavior $\beta = (\tau, \gamma, \mathbb{T})$, where $\tau = s_1\alpha_1s_2\alpha_2\dots$, and $\mathbb{T} = t_{s_1}t_{\alpha_1}t_{s_2}t_{\alpha_2}\dots$.

Definition 3 (Words and time sequences). We denote by $v_\tau = \varpi_1\varpi_2\dots = \mathcal{L}_i(\alpha_1)\mathcal{L}_i(\alpha_2)\dots \in (2^{\Sigma_i} \cup 2^{\mathcal{E}_i})^\omega$ the service set sequence associated with τ . The word w_τ produced by τ is the subsequence of the non-silent elements of v_τ ; $w_\tau = \varpi_{l_1}\varpi_{l_2}\dots \in (2^{\Sigma_i})^\omega$, such that $\varpi_1, \dots, \varpi_{l_1-1} = \mathcal{E}_i$, and for all $j \geq 1$, $\varpi_{l_j} \neq \mathcal{E}_i$ and $\varpi_{l_j+1}, \dots, \varpi_{l_{j+1}-1} = \mathcal{E}_i$. With a slight abuse of notation, we use $\mathbb{T}(v_\tau) = t_1t_2\dots = t_{\alpha_1}t_{\alpha_2}\dots$ to denote the service time sequence, i.e. the subsequence of \mathbb{T} when the services are provided. Furthermore, $\mathbb{T}(w_\tau) = t_{l_1}t_{l_2}\dots$ denotes the word time sequence, i.e. the subsequence of $\mathbb{T}(v_\tau)$ that corresponds to the time instances when the non-silent services are provided.

As in this work we are interested in infinite behaviors, we consider as *valid* traces only those producing infinite words.

Definition 4 (Service set at a time). Let τ be a trace of \mathcal{T}_i with $v_\tau = \varpi_1\varpi_2\dots$, and $\mathbb{T}(v_\tau) = t_1t_2\dots$. Given $t \in \mathbb{R}_0^+$, the service set $v_\tau(t) \in 2^{\Sigma_i} \cup 2^{\mathcal{E}_i}$ provided at time t is $v_\tau(t) = \varpi_j$ if $t = t_j$ for some $j \geq 1$; and $v_\tau(t) = \mathcal{E}_i$ otherwise.

Note that the above definition is designed in a way accommodating the general asynchronicity of the agents. Each formula ψ_i is interpreted locally, from the agent i 's point of view, based on the word w_{τ_i} it produces and on the services of the other agents $i' \in \mathcal{N}$ provided at the time instances $\mathbb{T}(w_{\tau_i})$. In other words, the agent i takes into consideration the other agents' services only at times, when i provides a non-silent service set (even an empty one) itself.

Definition 5 (Local LTL satisfaction). Let $\mathfrak{B} = \beta_1, \dots, \beta_N$ be a collection of behaviors, where $\beta_i = (\tau_i, \gamma_i, \mathbb{T}_i)$ for all $i \in \mathcal{N}$ and let $\mathbb{T}(w_{\tau_i}) = t_{i,l_1}t_{i,l_2}\dots$ be the word time sequence of agent i . The local word produced by \mathfrak{B} is $w_{\mathfrak{B}_i} = \omega_{i,l_1}\omega_{i,l_2}\dots$, where $\omega_{i,l_j} = (\bigcup_{i' \in \mathcal{N}} v_{\tau_{i'}}(t_{i,l_j})) \cap \Sigma$, for all $j \geq 1$. \mathfrak{B} locally satisfies the formula ψ_i for the agent i , $\mathfrak{B} \models_i \psi_i$, iff $w_{\mathfrak{B}_i} \models \psi_i$.

Note that even if $\psi_i = \psi_j$, it is possible that $\mathfrak{B} \models_i \psi_i$ for the agent i , while $\mathfrak{B} \not\models_j \psi_j$ for the agent j , since it might be the case that $w_{\mathfrak{B}_i} \neq w_{\mathfrak{B}_j}$.

C. Problem Statement

Problem 1. Consider a set of agents $\mathcal{N} = \{1, \dots, N\}$, and suppose that each agent $i \in \mathcal{N}$ is modeled as a tuple $\mathcal{M}_i = (\mathcal{T}_i = (S_i, s_{\text{init},i}, A_i, T_i, \Pi_i, L_i), \text{Sync}_i, \Sigma_i, \mathcal{L}_i)$, and assigned a task in the form of an LTL $_{\setminus X}$ formula ϕ_i over Π_i and ψ_i over $\Sigma = \bigcup_{i' \in \mathcal{N}} \Sigma_{i'}$. For each $i \in \mathcal{N}$ find a strategy,

i.e., (i) a trace $\tau_i = s_{i,1}\alpha_{i,1}s_{i,2}\alpha_{i,2}\dots$ of \mathcal{T}_i and (ii) a synchronization sequence γ_i over Sync_i with the property that the set of induced behaviors \mathbb{B} from Def. 2 is nonempty, and for all $\mathfrak{B} \in \mathbb{B}$ and all $i \in \mathcal{N}$, it holds that the trace τ_i satisfies ϕ_i and the word $w_{\mathfrak{B}_i}$ produced by \mathfrak{B} locally satisfies ψ_i for the agent i in terms of Def. 5.

As each LTL formula can be translated into a BA, from now on, we pose the problem with the motion specification of each agent i given as a BA $\mathcal{B}_i^\phi = (Q_i^\phi, q_{init,i}^\phi, \delta_i^\phi, 2^\Pi, F_i^\phi)$, and the task one as a BA $\mathcal{B}_i^\psi = (Q_i^\psi, q_{init,i}^\psi, \delta_i^\psi, 2^\Sigma, F_i^\psi)$. Given that $\tau_i = s_{i,1}\alpha_{i,1}s_{i,2}\alpha_{i,2}\dots$, the motion specification satisfaction is formulated as $L(s_{i,1})L(s_{i,2})\dots \in \text{Lang}(\mathcal{B}_i^\phi)$ and the local satisfaction of the task specification as $w_{\mathfrak{B}_i} \in \text{Lang}(\mathcal{B}_i^\psi)$.

IV. PROBLEM SOLUTION

Even though the agents' motion specifications are mutually independent, each of them is dependent on the agent's respective task specification, which is dependent on the task specifications of the other agents. As a result, the procedures of synthesizing the desired N strategies cannot be decentralized in an obvious way. However, one can quite easily obtain a centralized solution when viewing the problem as a synthesis of a single team strategy.

First of all, we realize that if there exists a solution to Problem 1, then there exists a solution in which all the agents synchronize after every single step. Informally, the reason is that each agent i is able to execute action $stay_i$, i.e., that $L_i(stay_i) = \mathcal{E}_i$, and that the motion specification ϕ_i is a formula of $\text{LTL}_{\setminus X}$, i.e., that the execution of $stay_i$ does not influence its satisfaction. The standard automata-based procedure for solving the LTL control strategy synthesis problem for deterministic systems can be then easily applied in multi-agent settings in a centralized way. 1) A synchronized TS $\mathcal{T}_{\mathcal{N}}$ with the set of states $S_{\mathcal{N}} = S_1 \times \dots \times S_N$ is built that represents all stepwise-synchronized traces of the agents in \mathcal{N} . 2) A BA $\mathcal{B}_{\mathcal{N}}$ is built for the formula $\bigwedge_{i \in \mathcal{N}} (\phi_i \wedge \psi_i)$ over the alphabet $2^{\bigcup_{i \in \mathcal{N}} \Pi_i \cup \Sigma}$. 3) A product automaton $\mathcal{P}_{\mathcal{N}}$ of $\mathcal{T}_{\mathcal{N}}$ and $\mathcal{B}_{\mathcal{N}}$ is constructed that captures only the agents' traces that are admissible by $\mathcal{T}_{\mathcal{N}}$ and result into behaviors that satisfy ϕ_1, \dots, ϕ_N , and locally satisfy ψ_1, \dots, ψ_N , for the agents $1, \dots, N$, respectively. The product $\mathcal{P}_{\mathcal{N}}$ is analyzed using standard graph algorithms in order to find its accepting run that projects onto the desired traces of \mathcal{T}_i , for all $i \in \mathcal{N}$. 4) The synthesized strategy for an agent $i \in \mathcal{N}$ is a trace $\tau_i = s_{i,1}\alpha_{i,1}s_{i,2}\alpha_{i,2}\dots$ of \mathcal{T}_i , and a synchronization sequence $\gamma_i = r_{i,1}r_{i,2}\dots$ where $r_{i,j} = \text{sync}_i\{\mathcal{N}\}$, for all $j \geq 1$. We notice that the amount of synchronization required by agent $i \in \mathcal{N}$ can be reduced by replacing $r_{i,j} = \text{sync}_i(\mathcal{N})$ with $\text{sync}_i(\{i\})$, if for each agent $i' \in \mathcal{N}$ it holds that $\mathcal{L}_{i'}(\alpha_{i',j}) = \mathcal{E}_{i'}$, since in such a case, $w_{\mathfrak{B}_i}$ remains unchanged. Furthermore, the $stay_i$ actions can be removed from τ_i without any harm, since they influence neither the satisfaction of ϕ_i , nor the word $w_{\mathfrak{B}_i}$ for any $i' \in \mathcal{N}$.

A. Our solution

A major drawback of the centralized solution is the state space explosion, which makes it practically unusable. We aim to decentralize the solution as much as possible. Namely, we aim to separate the synthesis of service plans yielding the local satisfaction of the task specifications from the syntheses of traces that guarantee the motion specifications. Our approach is to pre-compute possible traces and represent them efficiently, abstracting away the features that are not significant for the synthesis of action plans. This abstraction serves as a guidance for the action and synchronization planning, which, by construction, allows for finding a trace complying with both the synthesized action and synchronization plans and the motion specification.

1) *Preprocessing the motion specifications:* Consider for now an agent $i \in \mathcal{N}$ modeled as $\mathcal{M}_i = (\mathcal{T}_i, \text{Sync}_i, \Sigma_i, \mathcal{L}_i)$, where $\mathcal{T}_i = (S_i, s_{init,i}, A_i, T_i, \Pi_i, L_i)$ and its motion specification BA \mathcal{B}_i^ϕ . We slightly modify the classical construction of a product automaton of \mathcal{T}_i and \mathcal{B}_i^ϕ to obtain a BA that represents the traces of \mathcal{T}_i accepted by \mathcal{B}_i^ϕ and furthermore explicitly captures the services provided along the trace.

Definition 6 (Motion product). *The motion product of a TS $\mathcal{T}_i = (S_i, s_{init,i}, A_i, T_i, \Pi_i, L_i)$, and a BA $\mathcal{B}_i^\phi = (Q_i^\phi, q_{init,i}^\phi, \delta_i^\phi, 2^{\Pi_i}, F_i^\phi)$ is a BA $\mathcal{P}_i = (Q_i, q_{init,i}, \delta_i, 2^{\Sigma_i} \cup \{\mathcal{E}_i\}, F_i)$, where $Q_i = S_i \times Q_i^\phi$; $q_{init,i} = (s_{init,i}, q_{init,i}^\phi)$; $((s, q), \mathcal{L}_i(\alpha), (s', q')) \in \delta_i$ if and only if $(s, \alpha, s') \in T_i$, and $(q, L_i(s), q') \in \delta_i^\phi$; and $F_i = \{(s, q) \mid q \in F_i^\phi\}$.*

Intuitively, we propose to keep only the significant states that have an outgoing transition labeled with $\mathcal{L}_i(\alpha) \neq \mathcal{E}_i$ and replace exact sequences of transitions in \mathcal{P}_i with single transitions representing reachability. While doing so, we have to take into account whether the removed state is accepting or not to correctly preserve the accepting condition of \mathcal{P}_i .

Definition 7 (Insignificant states in \mathcal{P}_i). *A state p of the motion product \mathcal{P}_i is significant if it is (i) the initial state $p = q_{init,i}$, or (ii) if there exists a transition $(p, \sigma, p') \in \delta_i$, such that $\sigma \neq \mathcal{E}_i$; and insignificant otherwise.*

A reduced motion product $\ddot{\mathcal{P}}_i$ is built from \mathcal{P}_i according to Alg. 1. First, we remove all insignificant non-accepting states and their incoming and outgoing transitions (lines 3, 4). We replace each state with a set of transitions leading directly from the state's predecessors to its successors, i.e., we concatenate the incoming and the outgoing transitions (line 5). The labels of the new transitions differ: if both labels of the concatenated incoming and outgoing transition are \mathcal{E}_i , then the new label will be \mathcal{E}_i to indicate that the transition represents a sequence of actions that are not interesting with respect to the local satisfaction of task specifications. On the other hand, if the label σ of the incoming transition belongs to 2^{Σ_i} , we use the action σ as the label for the new transition. Each path between two significant states in \mathcal{P}_i maps onto a path between the same states in $\ddot{\mathcal{P}}_i$ and the sequences of non-silent services read on the labels of the transitions of the two paths are equal; and vice versa. Second, we handle the insignificant accepting states (line 7) similarly to the non-

accepting ones, however, we do not remove the states whose predecessors include a significant state in order to preserve the accepting condition. There is a correspondence between the infinite runs of \mathcal{P}_i and the infinite runs of $\bar{\mathcal{P}}_i$: for each run of \mathcal{P}_i there exists a run of $\bar{\mathcal{P}}_i$, such that the states of the latter one are a subsequence of the states of the former one, the sequences of non-silent services read on the labels of the transitions of the two runs are equal, and that the latter one is accepting if and only if the former one is accepting; and vice versa. This correspondence will allow us to reconstruct a desired run of \mathcal{P}_i from a run of $\bar{\mathcal{P}}_i$, as we will show in Sec. IV-A.4. From $\bar{\mathcal{P}}_i$, we further remove all states from which none of the states in \bar{F}_i is reachable, and we can keep only one copy of duplicate states that have analogous incoming and outgoing edges (omitted from Alg. 1).

Algorithm 1 Reduction of the motion product

Input: motion product $\mathcal{P}_i = (Q_i, q_{init,i}, \delta_i, 2^{\Sigma_i} \cup 2^{\mathcal{E}_i}, F_i)$
Output: reduced BA $\bar{\mathcal{P}}_i = (\bar{Q}_i, \bar{q}_{init,i}, \bar{\delta}_i, 2^{\Sigma_i} \cup 2^{\mathcal{E}_i}, \bar{F}_i)$

- 1: initialize $\bar{\mathcal{P}}_i := \mathcal{P}_i$
- 2: **for all** insignificant states $p \in Q_i \setminus F_i$ **do**
- 3: $\bar{Q}_i := \bar{Q}_i \setminus \{p\}$
- 4: $\bar{\delta}_i := \delta_i \setminus \{(p', \sigma, p), (p, \sigma, p') \mid p' \in \bar{Q}_i, \sigma \in 2^{\Sigma_i} \cup 2^{\mathcal{E}_i}\}$
- 5: $\bar{\delta}_i := \bar{\delta}_i \cup \{(p', \sigma, p'') \mid p', p'' \in \bar{Q}_i, \sigma \in 2^{\Sigma_i} \cup 2^{\mathcal{E}_i}, (p', \sigma, p), (p, \mathcal{E}_i, p'') \in \delta_i\}$
- 6: **end for**
- 7: **for all** insignificant states $p \in F_i$, such that all predecessors of p in $\bar{\mathcal{P}}_i$ are insignificant **do**
- 8: **if** $(p, \mathcal{E}_i, p) \in \bar{\delta}_i$ **then**
- 9: $\bar{\delta}_i := \bar{\delta}_i \cup \{(p', \mathcal{E}_i, p') \mid (p', \mathcal{E}_i, p) \in \bar{\delta}_i\}$
- 10: **end if**
- 11: $\bar{Q}_i := \bar{Q}_i \setminus \{p\}$
- 12: $\bar{\delta}_i := \bar{\delta}_i \setminus \{(p', \sigma, p), (p, \sigma, p') \mid p' \in \bar{Q}_i, \sigma \in 2^{\Sigma_i} \cup 2^{\mathcal{E}_i}\}$
- 13: $\bar{\delta}_i := \bar{\delta}_i \cup \{(p', \mathcal{E}_i, p'') \mid p', p'' \in \bar{Q}_i, (p', \mathcal{E}_i, p), (p, \mathcal{E}_i, p'') \in \bar{\delta}_i\}$
- 14: **end for**

2) *Preprocessing the task specifications:* Next, we build a local task and motion product automaton for each agent i separately, to capture the admissible traces of i that comply both with its motion and task specification. At this stage, the other agents' collaboration capabilities are not yet included.

Definition 8 (Task and motion product). *The task and motion product of a reduced motion product automaton $\bar{\mathcal{P}}_i = (\bar{Q}_i, \bar{q}_{init,i}, \bar{\delta}_i, 2^{\Sigma_i} \cup 2^{\mathcal{E}_i}, \bar{F}_i)$, and the task specification BA $\mathcal{B}_i^\psi = (Q_i^\psi, q_{init,i}^\psi, \delta_i^\psi, 2^{\Sigma_i}, F_i^\psi)$ is a BA $\bar{\mathcal{P}}_i = (\bar{Q}_i, \bar{q}_{init,i}, \bar{\delta}_i, 2^{\Sigma_i} \cup 2^{\mathcal{E}_i}, \bar{F}_i)$, where $\bar{Q}_i = \bar{Q}_i \times Q_i^\psi \times \{1, 2, 3\}$; $\bar{q}_{init,i} = (q_{init,i}, q_{init,i}^\psi, 1)$; $\bar{F}_i = \{(q_1, q_2, 2) \mid q_2 \in F_i^\psi\}$; and $((q_1, q_2, j), \sigma, (q'_1, q'_2, j')) \in \bar{\delta}_i$ iff*

- $\sigma = \mathcal{E}_i$, $(q_1, \mathcal{E}_i, q'_1) \in \bar{\delta}_i$, and $q_2 = q'_2$; or
 - $\sigma \in 2^{\Sigma_i}$, $(q_1, \sigma \cap 2^{\Sigma_i}, q'_1) \in \bar{\delta}_i$, and $(q_2, \sigma, q'_2) \in \delta_i^\psi$,
- and $j' = 2$, if $j = 1$ and $q'_1 \in \bar{F}_i$, $j' = 3$ if $j = 2$ and $q'_2 \in \bar{F}_i^\psi$, $j' = 1$ if $j = 3$, and $j = j'$ otherwise.

The above construction leverages ideas from building a BA that accepts a language intersection of multiple given BAs. The accepting runs of $\bar{\mathcal{P}}_i$ map onto accepting runs of $\bar{\mathcal{P}}_i$, and hence to the traces of \mathcal{T}_i satisfying the motion specification ϕ_i , and onto accepting runs of \mathcal{B}_i^ψ ; and vice versa. Clearly, some of the transitions of $\bar{\mathcal{P}}_i$ require a

collaboration of some other agents, while some of them do not. In order to further reduce the size of $\bar{\mathcal{P}}_i$, we introduce a mapping $Dep_i : \bar{\delta}_i \rightarrow 2^{\mathcal{N}}$ to indicate for each transition $t = (p, \sigma, p') \in \bar{\delta}_i$ the set of agents $Dep_i(t) \subseteq \mathcal{N}$ whose collaboration is required. We formalize Dep_i through the notion of assisting services. Moreover, for each agent i , we find a subset of its services that affect the local satisfaction of a task specification of another agent through the notion of globally assisting services.

Definition 9 (Assisting services). *Suppose that $i \neq i'$. A service $\rho \in \Sigma_{i'}$ is not assisting on a transition $(p, \sigma, p') \in \bar{\delta}_i$ of $\bar{\mathcal{P}}_i$ if and only if it holds that $(p, \sigma \cup \{\rho\}, p') \in \bar{\delta}_i \iff (p, \sigma \setminus \{\rho\}, p') \in \bar{\delta}_i$; it is assisting on (p, σ, p') otherwise. $Dep_i(t) = \{i\} \cup \{i' \mid i' \neq i \text{ and } \exists \rho \in \Sigma_{i'} \text{ assisting on } t\}$.*

Definition 10 (Globally assisting services). *A service $\rho \in \Sigma_{i'}$ of agent $i' \in \mathcal{N}$ is globally assisting if there exists $i \in \mathcal{N}$, $i \neq i'$, and a transition $t \in \bar{\delta}_i$, such that ρ is assisting on t .*

Now, similar ideas as in Alg. 1 can be used to reduce the size of each action and motion product $\bar{\mathcal{P}}_i$, with the following, altered definition of significant states:

Definition 11 (Insignificant states in $\bar{\mathcal{P}}_i$). *A state p of the task and motion product $\bar{\mathcal{P}}_i$ is significant if it is either (i) the initial state $p = \bar{q}_{init,i}$, or (ii) if there exists a transition $(p, \sigma, p') \in \bar{\delta}_i$, such that there exists a globally assisting service $\rho \in \sigma \cap \Sigma_i$, or (iii) if there exists a transition $(p, \sigma, p') \in \bar{\delta}_i$, such that $Dep_i(p, \sigma, p') \neq \{i\}$; the state p is insignificant otherwise.*

Informally, the condition (ii) states that at a significant state, the agent i has the ability to assist other agents, whereas the condition (iii) states that at a significant state, the assistance of the other agents influences the agent i 's own task specification local satisfaction. With a slight abuse of notation, we replace all labels of all outgoing transitions of insignificant states of $\bar{\mathcal{P}}_i$ with \mathcal{E}_i . This time \mathcal{E}_i represents a service set that is completely independent. We can now directly apply Alg. 1 to remove the insignificant states from $\bar{\mathcal{P}}_i$, with the exception that 2^{Σ_i} is at all places replaced with 2^{Σ_i} . The resulting automaton is $\hat{\mathcal{P}}_i = (\hat{Q}_i, \hat{q}_{init,i}, \hat{\delta}_i, 2^{\Sigma_i} \cup 2^{\mathcal{E}_i}, \hat{F}_i)$. Similarly as before, there is a correspondence between the infinite runs of $\bar{\mathcal{P}}_i$ and the infinite runs of $\hat{\mathcal{P}}_i$: for each run of $\bar{\mathcal{P}}_i$ there exists a run of $\hat{\mathcal{P}}_i$, such that the states of the latter one are a subsequence of the states of the former one, the sequences of services read on the labels of the transitions leading from the significant states of the two runs are equal, and that the latter one is accepting if and only if the former one is accepting; and vice versa. The number of states of $\hat{\mathcal{P}}_i$ is at most twice the number of significant states in $\bar{\mathcal{P}}_i$.

3) *Global product:* From the reduced task and motion product automata $\hat{\mathcal{P}}_1, \dots, \hat{\mathcal{P}}_N$, we build a single one, that finally represents the inter-agent collaborations. The construction is a twist to the well-known construction of BA intersection. We associate the transitions of the BA with the subsets of agents that are required to make this transition collaboratively, i.e., with the subsets of agents that need to synchronize prior this transition.

Definition 12 (Global product). *The global product of the reduced task and motion product automata $\hat{\mathcal{P}}_1, \dots, \hat{\mathcal{P}}_N$, where $\hat{\mathcal{P}}_i = (\hat{Q}_i, \hat{q}_{init,i}, \hat{\delta}_i, 2^{\Sigma} \cup 2^{\mathcal{E}_i}, \hat{F}_i)$, for all $i \in \mathcal{N}$, is a BA $\mathcal{P} = (Q, q_{init}, \delta, 2^{\Sigma} \cup 2^{\mathcal{E}=\{\mathcal{E}_i | i \in \mathcal{N}\}}, F)$ with a mapping $Dep : \delta \rightarrow 2^{\mathcal{N}}$, where $Q = \hat{Q}_1 \times \dots \times \hat{Q}_N \times \{1, \dots, N+1\}$; $q_{init} = (\hat{q}_{init,1}, \dots, \hat{q}_{init,N}, 1)$; $\hat{F}_i = \{(q_1, \dots, q_N, N) \mid q_N \in \hat{F}_N\}$; and $t = ((q_1, \dots, q_N, j), \sigma, (q'_1, \dots, q'_N, j')) \in \delta$ iff either*

- $\exists i \in \mathcal{N}$, such that $\sigma = \mathcal{E}_i$, $(q_i, \mathcal{E}_i, q'_i) \in \hat{\delta}_i$, $q_{i'} = q'_{i'}$, for all $i' \neq i$, and $j' = j + 1$ if $j = i$ and $q'_i \in \hat{F}_i$, $j' = 1$ if $j = N + 1$, and $j' = j$ otherwise. Then we set $Dep(t) = \{i\}$; or
- $\sigma \in 2^{\Sigma}$, and $\exists \mathcal{I} \subseteq \mathcal{N}$, such that for all $i \in \mathcal{I}$ it holds that $(q_i, \sigma, q'_i) \in \hat{\delta}_i$ while for all $i \notin \mathcal{I}$ it holds that $q_i = q'_i$. Moreover, $\bigcup_{i \in \mathcal{I}} Dep_i(q_i, \sigma, q'_i) \subseteq \mathcal{I}$, and $j' = j + 1$ if $j \in \mathcal{I}$ and $q'_j \in \hat{F}_j$, $j' = 1$ if $j = N + 1$, and $j' = j$ otherwise. Then we set $Dep(t) = \mathcal{I}$.

Each accepting run of the global product \mathcal{P} maps directly on the accepting runs of the reduced task and motion product automata and vice versa, for each collection of accepting runs of the reduced task and motion product automata, there exists an accepting run of the global product \mathcal{P} . Note that by this construction, deadlocks are avoided.

4) *Synthesis of the strategies*: The final step of our solution is the generation of the strategy in \mathcal{P} and its mapping onto a trace τ_i of \mathcal{T}_i and a synchronization sequence γ_i over $Sync_i$, for all $i \in \mathcal{N}$. Using standard graph algorithms (see, e.g., [1]), we find an accepting run $q_1 q_2 \dots$ over a word $\sigma_1 \sigma_2 \dots$ in \mathcal{P} , where $q_j = (\hat{q}_{1,j}, \dots, \hat{q}_{N,j}, k)$, for all $j \geq 1$. For each agent $i \in \mathcal{N}$:

- (i) Consider the sequence $\hat{q}_{i,1} \hat{q}_{i,2} \dots$, that is obtained by the projection of the accepting run $q_1 q_2 \dots$ onto the states of $\hat{\mathcal{P}}_i$. Let $\iota_1 \iota_2 \dots$ be the subsequence of all indexes, such that $i \in Dep(q_{\iota_j}, \sigma_{\iota_j}, q_{\iota_j+1})$. Then $\hat{q}_{i,\iota_1} \hat{q}_{i,\iota_2} \dots$ is an accepting run of $\hat{\mathcal{P}}_i$ over the word $\sigma_{\iota_1} \sigma_{\iota_2} \dots$ and $\hat{\gamma}_i = \hat{r}_{i,1} \hat{r}_{i,2} \dots = Dep(q_{\iota_1} \sigma_{\iota_1} q_{\iota_1+1}) Dep(q_{\iota_2} \sigma_{\iota_2} q_{\iota_2+1}) \dots$ is a synchronization sequence.
- (ii) From the construction of $\hat{\mathcal{P}}_i$, we know that there exists an accepting run $\bar{q}_{i,1} \bar{q}_{i,2} \dots$ in $\bar{\mathcal{P}}_i$ over a word $\bar{\sigma}_{i,1} \bar{\sigma}_{i,2} \dots$, such that $\bar{q}_{i,\iota_1} \bar{q}_{i,\iota_2} \dots$ is a subsequence $\bar{q}_{i,\ell_1} \bar{q}_{i,\ell_2} \dots$ of $\bar{q}_{i,1} \bar{q}_{i,2} \dots$ and $\sigma_{\iota_1} \sigma_{\iota_2} \dots$ is the corresponding subsequence $\bar{\sigma}_{i,\ell_1} \bar{\sigma}_{i,\ell_2} \dots$ of $\bar{\sigma}_{i,1} \bar{\sigma}_{i,2} \dots$. The synchronization sequence $\bar{\gamma}_i = \bar{r}_{i,1} \bar{r}_{i,2} \dots$ is constructed by setting $\bar{r}_{i,\ell_j} = \hat{r}_{i,\iota_j}$, for all $j \geq 1$, and $\bar{r}_{i,j} = sync(\{i\})$ otherwise.
- (iii) The mapping from the task and motion product $\bar{\mathcal{P}}_i$ onto $\bar{\mathcal{P}}_i$ is processed as follows: Suppose that $\bar{q}_{i,j} = (\bar{q}_{i,j}, q_{i,j}^\psi, k)$, for all $j \geq 1$. The accepting run $\bar{q}_{i,\ell_1} \bar{q}_{i,\ell_2} \dots$ then maps onto the accepting run $\bar{q}_{i,\ell_1} \bar{q}_{i,\ell_2} \dots$ over $\bar{\sigma}_{i,\ell_1} \bar{\sigma}_{i,\ell_2} \dots$, where $\bar{\sigma}_{i,j} = \mathcal{E}_i$ if $\bar{\sigma}_{i,j} = \mathcal{E}_i$, and $\bar{\sigma}_{i,j} = \bar{\sigma}_{i,j} \cap 2^{\Sigma_i}$ otherwise. Naturally, $\bar{\gamma}_i = \bar{\gamma}_i$.
- (iv) The accepting run $q_{i,1} q_{i,2} \dots$ of the motion product \mathcal{P}_i over a word $\sigma_{i,1} \sigma_{i,2} \dots$ and a synchronization sequence γ_i is obtained from the reduced motion product $\bar{\mathcal{P}}_i$ by an analogous construction as in the case of mapping from the reduced task and motion product $\hat{\mathcal{P}}_i$ onto the

task and motion product $\bar{\mathcal{P}}_i$.

- (v) Finally, an accepting run $q_{i,1} q_{i,2} \dots$ of \mathcal{P}_i maps onto a trace τ_i of \mathcal{T}_i . The desired synchronization sequence is γ_i .

B. Discussion

Strategy execution and event-based recomputation: In the above described solution, the strategy for each agent is constructed in an offline manner, and executed independently. However, the solution can be modified to an event-triggered one, where synchronization is triggered by need and strategies are periodically recomputed in order to adapt to different execution speed of different agents and to handle various optimization criteria. The details of this solution is one of our future topics of interest. *Mutually dependency*: We supposed that all of the considered agents are mutually dependent through their task specifications, either directly or transitively. If this is not the case, the team of agents can be partitioned into dependency classes as suggested in [5]. *Complexity*: In the worst case, our solution meets the complexity of the centralized solution. However, this is often not the case. Since the size of the global product is highly dependent on the number of globally assisting services available in the agents' workspace, our solution is particularly suitable for systems with complex motion capabilities, sparsely distributed services of interest, and occasional needs for collaboration. Its benefits are demonstrated in Sec. V.

V. ILLUSTRATIVE EXAMPLE

An illustrative example with 3 robots is given in Fig. 1. Their workspace is partitioned into 100 cells, and a robot's state is defined by the cell it is present at. Agent 1 (green) is a ground vehicle and has to avoid the obstacles (in gray), while agents 2 (orange) and 3 (purple) are UAVs and their workspace is obstacle-free except for walls (thick). The cells are each labeled with an atomic proposition from $\{R1, R2, R3, R4\}$ indicating the room to which the cell belongs. Services are available in cell marked with stripes. Agent 1 can *load* (green horizontal stripe) and *unload* (green vertical). Agent 2 can *help* (orange horizontal) or *inform* (orange vertical). Agent 3 can *assist* (purple horizontal). The motion formulas are: For agent 1 to avoid room R1, $\phi_1 = G \neg R1$, for agent 2 to avoid R2, $\phi_2 = G \neg R2$, and for agent 3 to survey R1 and R2, $\phi_3 = GF R1 \wedge GF R2$. The task formulas are: For agent 1 to periodically load collaboratively with agents 2 and 3 and then to unload, collaboratively with agent 2 or 3: $\psi_1 = load \wedge help \wedge assist \wedge G (load \Rightarrow X (unload \wedge (help \vee assist))) \wedge G (unload \Rightarrow X (load \wedge help \wedge assist))$, for agent 2 to periodically service *inform*, $\psi_2 = GF inform$, for agent 3 $\psi_3 = assist \vee \neg assist$, i.e., agent 3 is not assigned a specific task. An example of as collection of the desired trajectories and synchronizations is given in Fig. 1.

In classical centralized planning, only the synchronous product of the TSs of the three agents would have ≈ 100 000 states. In our case, the BAs for the given formulas had 2,2,3,2,2,1 states, respectively, and hence their language intersection BAs reaches $2 * 2 * 3 * 2 * 2 * 1 * 7 \approx 330$ states.




Fig. 1: Examples of desired trajectories of 3 robots in a shared workspace. The synchronization events at the moments of providing nonsilent services are depicted with stars. First, all agents synchronize in the bottom left corner and provide *load*, *help*, *assist*; then agent 2 provides *inform* and then agents 1 and 3 provide *unload*, *assist*, or vice versa; finally, all agents synchronize in the bottom left corner, and start periodically executing the above.

The overall product automaton has then up to ~ 30 millions states. In contrast, in our solution, the individual task and motion products $\hat{\mathcal{P}}_1, \hat{\mathcal{P}}_2, \hat{\mathcal{P}}_3$ after all reductions have 27, 17, and 8 states, respectively. Hence the largest structure dealt with during the overall procedure, i.e., the product automaton \mathcal{P} has only $27 * 17 * 8 * 4 \sim 15000$ states.

REFERENCES

- [1] C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
- [2] A. Bhatia, M. R. Maly, L. E. Kavraki, and M. Y. Vardi. Motion planning with complex goals. *Robotics Automation Magazine, IEEE*, 18(3):55–64, 2011.
- [3] Y. Chen, X. C. Ding, A. Stefanescu, and C. Belta. Formal approach to the deployment of distributed robotic teams. *IEEE Transactions on Robotics*, 28(1):158–171, 2012.
- [4] I. Filippidis, D. V. Dimarogonas, and K. J. Kyriakopoulos. Decentralized multi-agent control from local ltl specifications. In *Proceedings of the IEEE Conference on Decision and Control*, pages 6235–6240, 2012.
- [5] M. Guo and D. V. Dimarogonas. Multi-agent plan reconfiguration under local LTL specifications. *International Journal of Robotics Research*, pages 218–235, 2015.
- [6] M. Guo, J. Tumova, and D. V. Dimarogonas. Cooperative decentralized multi-agent control under local LTL tasks and connectivity constraints. In *Proceedings of the IEEE Conference on Decision and Control*, pages 75–80, 2014.
- [7] G. Jing, C. Finucane, V. Raman, and H. Kress-Gazit. Correct high-level robot control from structured english. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3543–3544, 2012.
- [8] S. Karaman and E. Frazzoli. Vehicle routing with temporal logic specifications: Applications to multi-UAV mission planning. *International Journal of Robust and Nonlinear Control*, 21:1372–1395, 2011.
- [9] M. Kloetzer and C. Belta. A fully automated framework for control of linear systems from temporal logic specifications. *IEEE Transactions on Automatic Control*, 53(1):287–297, 2008.
- [10] M. Kloetzer, X. C. Ding, and C. Belta. Multi-robot deployment from LTL specifications with reduced communication. In *Proceedings of the IEEE Conference on Decision and Control and European Control Conference*, pages 4867–4872, 2011.
- [11] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas. Temporal logic-based reactive mission and motion planning. *IEEE Transactions on Robotics*, 25(6):1370–1381, 2009.
- [12] S. G. Loizou and K. J. Kyriakopoulos. Automated planning of motion tasks for multi-robot systems. In *Proceedings of the IEEE Conference on Decision and Control*, volume 44, 2005.
- [13] J. Tumova and D. Dimarogonas. A receding horizon approach to multi-agent planning from local LTL specifications. In *Proceedings of the American Control Conference*, pages 1775–1780, 2014.

- [14] J. Tumova and D. Dimarogonas. Multi-agent planning under local LTL specifications and event-based synchronization. 2015. Submitted.
- [15] A. Ulusoy, S. L. Smith, X. C. Ding, C. Belta, and D. Rus. Optimality and robustness in multi-robot path planning with temporal logic constraints. *International Journal of Robotics Research*, 32(8):889–911, 2013.
- [16] C. Wiltse, F. A. Ramponi, and J. Lygeros. Synthesis of an asynchronous communication protocol for search and rescue robots. In *Proceedings of the European Control Conference*, pages 1256–1261, 2013.
- [17] T. Wongpiromsarn, U. Topcu, and R. M. Murray. Receding Horizon Control for Temporal Logic Specifications. In *Hybrid systems: Computation and Control (HSCC)*, pages 101–110, 2010.